Autonomous Warplanes

NASA Rovers Lead the Way

Michael R. Schroer
Major, Air National Guard

AIR UNIVERSITY
AIR COMMAND AND STAFF COLLEGE
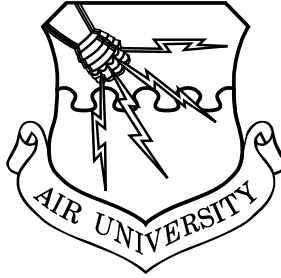
Autonomous Warplanes

NASA Rovers Lead the Way

Michael R. Schroer
Major, Air National Guard

Wright Flyer No. 54

Air University Press
Air Force Research Institute
Maxwell Air Force Base, Alabama

Published by Air University Press in April 2016

AFRI AU PRESS
AIR FORCE RESEARCH INSTITUTE

This Wright Flyer Paper and others in the series are available electronically at the AU Press website: http://aupress.au.af.mil.
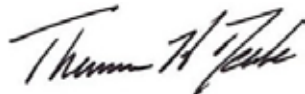
Contents

Foreword

It is my great pleasure to present another issue of The Wright Flyer Papers. Through this series, Air Command and Staff College (ACSC) presents a sampling of exemplary research produced by our residence and distance-learning students. This series has long showcased the kind of visionary thinking that drove the aspirations and activities of the earliest aviation pioneers. This year's selection of essays admirably extends that tradition. As the series title indicates, these papers aim to present cutting-edge, actionable knowledge— research that addresses some of the most complex security and defense challenges facing us today.

Recently, The Wright Flyer Papers transitioned to an exclusively electronic publication format. It is our hope that our migration from print editions to an electronic-only format will fire even greater intellectual debate among Airmen and fellow members of the profession of arms as the series reaches a growing global audience. By publishing these papers via the Air University Press website, ACSC hopes not only to reach more readers, but also to support Air Force–wide efforts to conserve resources. In this spirit, we invite you to peruse past and current issues of The Wright Flyer Papers at http://aupress.maxwell.af.mil/papers_all.asp?cat=wright.

Thank you for supporting The Wright Flyer Papers and our efforts to disseminate outstanding ACSC student research for the benefit of our Air Force and war fighters everywhere. We trust that what follows will stimulate thinking, invite debate, and further encourage today's air, space, and cyber war fighters in their continuing search for innovative and improved ways to defend our nation and way of life.

THOMAS H. DEALE
Brigadier General, USAF
Commandant

About the Author

After receiving a commission from the US Air Force Academy in 2000, Maj Michael Schroer earned his pilot wings at Vance AFB, Oklahoma, and has enjoyed a varied flying career in the subsequent 15 years. At Andrews AFB, Maryland, the demands of distinguished visitor transportation between most airports across the continent proved an excellent further education in aviation. Piloting a business jet on a weeklong, 11-hop trek across the North Atlantic and Europe to Qatar was a highlight of the assignment. At Offutt AFB, Nebraska, Major Schroer flew manned reconnaissance aircraft, primarily the RC-135S Cobra Ball and the related RC-135U Combat Sent.

After a dozen years of active service, Major Schroer secured a position in the Air National Guard with the famous "Happy Hooligans" flying the MQ-1 Predator. The combination of demanding mission sets and the excitement of being part of a fundamental shift in the trajectory of aviation has proven ample compensation for the loss of cockpit windows. Remotely piloted aircraft (RPA) are a natural fit for his longtime interests in computers, technology, and cameras. Perhaps most importantly, remote piloting offers him the unique ability to contribute to Air Force operations without the penalty of frequent and lengthy separation from family. Major Schroer looks forward to playing a part in the continued rapid evolution of the RPA from a technology demonstration to an indispensible tool of the modern world.

Acknowledgments

The very mention of the words "autonomous" and "lethal" in the same sentence is bound to provoke an emotional response. The use of autonomous aircraft in warfare is a controversial topic that is rapidly rising to the surface of public consciousness. In July 2015 more than 1,000 artificial intelligence researchers penned an open letter calling for a ban on the development of autonomous weapons. Signatories include noteworthy thinkers and industry leaders, including such recognizable personalities as physicist Stephen Hawking, Tesla chief executive officer Elon Musk, Apple cofounder Steve Wozniak, philosopher Daniel Dennett, and many others. These leaders raise valid concerns about the eventual evolution of the technologies discussed in this paper, likening their importance to the invention of gunpowder and nuclear weapons. I hope that the following discussion provides a useful perspective on some of the relevant issues at play in this complex subject.

Abstract

The capabilities and use of remotely piloted aircraft (RPA) are growing at a remarkable pace, quickly becoming indispensable parts of military operations and assets greatly valued by commanders in the field. With advancements in technology, RPAs are becoming more autonomous or capable of performance with less direct human control. Research conducted by the National Aeronautics and Space Administration (NASA) offers useful lessons for the development of future military RPAs that may allow the United States to maintain its current status as a world leader in the rapidly evolving RPA field.

This paper employs the problem/solution methodology to identify those aspects of NASA's rover autonomy research that may apply to military RPA development. Specifically, the concept of modular design and the rovers' feature detection, planning, scheduling, and prioritization subsystems offer solutions to current RPA challenges. Additionally, NASA's successful implementation of adjustable autonomy demonstrates an important transitional model that will be critical for the incremental transition from remotely piloted to autonomous warplanes. The pursuit of fully autonomous warplanes will persist due to the numerous advantages offered by autonomous operation and competitive pressures in the international arena. This paper recommends that the US Air Force mimic the design principles NASA has proven to be effective and, in particular, adhere to a modular approach in its procurement of future RPA systems.

Introduction

Military use of remotely piloted aircraft (RPA) has grown dramatically in recent years. An RPA is an aircraft flown by a human pilot not physically present in the aircraft. The pilot sends control inputs via a data link that can, through the use of satellite communications, reach across the globe from pilot to aircraft. In turn, the aircraft sends information about its performance and collected data back to the pilot along the same link. These systems offer many benefits, including reduced or eliminated risk to human life, increased efficiency, dramatically improved time on station, and lower cost. However, two key factors can influence the performance of current systems: the requirement for direct—albeit remote—human control due to the inherently slow process of decision making and the physical delay, or latency, of satellite communications. The amount of information that can be transmitted around the world via satellite, or bandwidth, is a finite resource because of the laws of physics and the hardware constraints that are ultimately tied to costs.

The practical implication for the current RPA fleet is that the amount and quality of video and telemetry a system can return to its operator are limited. This reality runs contrary to a seemingly insatiable appetite of combatant commanders and intelligence analysts for more video feeds, from more locations, and with better resolution. The current architecture creates military vulnerabilities aside from the limitations imposed by bandwidth and latency. All of this communication is routed through a small number of ground stations and satellites susceptible to degradation from space weather and atmospheric weather phenomena, hardware failure, jamming (both intentional and inadvertent), and physical attack. At present, a single thunderstorm can disable a large and growing percentage of the US Air Force's available intelligence, surveillance, and reconnaissance (ISR) and combat airpower. Imagine the pilots of half the combat aircraft in-theater simultaneously developing narcolepsy; the severity of such an exigency is clear.

The introduction of ever-greater levels of autonomy through replacing direct human control with onboard decision making mitigates the problems of bandwidth scarcity, communication latency, and vulnerability by eliminating the bulk of that communication. Additionally, autonomy offers the promise of advantages in decreased decision time, greater effectiveness, and reduced need for human operators. Moreover, the rapid decision making and responsiveness of automated airborne systems could allow these platforms to operate reliably in contested or denied territory—a considerable drawback of current systems.

The technical demands of developing sufficiently autonomous warplanes are daunting but not entirely new; the National Aeronautics and Space Administration (NASA) has dealt with similar issues for decades in its robotic exploration of space. While the limited bandwidth and latency of communications around the world are concerns, the constraints inherent to controlling a mobile robot in another planet's harsh environment are far more significant. As technological advances have increased the range, speed, power, and instrument resolution of these rovers, NASA has responded by endowing them with an increasing ability to act independently. This initiative is organized under the onboard autonomous science investigation system (OASIS) program. OASIS incorporates a modular approach to developing autonomy, pursuing advances in the feature recognition, prioritization, scheduling, and planning systems as components of an integrated whole. Applying NASA's research to military machines could allow the United States to take advantage of decades of development effort and maintain the technological advantage it currently enjoys in this rapidly evolving area of competition.

How does NASA's interplanetary rover research offer solutions to the challenges of burgeoning autonomy in RPAs? The modular approach to machine learning, scheduling, planning, and stereoscopic vision systems developed for NASA's OASIS program offers a useful model for the development of autonomy in future USAF RPA systems. This study explores the potential for applying NASA's research and solutions related to the rover to RPA systems.

Just as communications latency and inadequate bandwidth impose limitations on current RPA operations, these factors became an impediment to NASA's pursuit of scientific goals—particularly when exacerbated by steadily improved rover performance. NASA's solution was to avoid the communication problems altogether by endowing its machines with increasing degrees of autonomy, thereby making them somewhat self-sufficient. Despite the different environments, distances, and purposes, military RPA systems are currently limited in many similar ways by the reliance on direct human control and the enabling communications. Greater autonomy in these systems would address those limitations and lead to meaningful improvements, including reliability, responsiveness, lethality, efficiency, and reduced human workload. These same principles and modular approach pursued by NASA's rover scientists can be used to develop future RPA systems.

This study uses the problem/solution methodology to explore some of the obstacles to increasing automation in military RPAs. It discusses the phenomenon of automation in its historical perspective and then addresses the problems associated with current RPA operations. NASA's

specific response to similar challenges in the space exploration domain is examined, and the applicability of these solutions to military RPAs is analyzed. Finally, this study provides recommendations for future development of military RPAs and addresses potential pitfalls to be avoided.

Background

Development of Automated Machines

Throughout history, humanity has sought ways to multiply the benefits of its labor; perhaps the most powerful example is the development of automation. Machines that perform tasks previously requiring human attention and labor can be leveraged to reduce human workload or increase individual effectiveness. Recent technological advances have approached the logical end point of this evolution—the development of autonomous machines capable of performing their intended function with little or no human oversight for substantial periods of time. Such machines offer obvious advantages, particularly in situations where direct human control is either impossible or undesirable.

In this study, "autonomy" reflects the ability of an electronic system to assess the external environment and make decisions concerning the system's interaction with that environment in light of certain goals, all without direct human control during the system's operation. The concept of autonomy is distinct from the larger concept of automation. Both refer to the completion of tasks formerly performed by humans, but automation refers to tasks done at specific human command and without any awareness of the local environment or goal-oriented decision making based upon that awareness. Factory robots are a prime example of automated machines that are not generally autonomous.

For example, a hypothetical autonomous aircraft could be given a high-level task, such as assessing a remote village's population. Using onboard sensors and linked data sources, this autonomous machine would determine the most appropriate path to its destination, taking terrain, weather, airspace boundaries, and other traffic into account. It would then coordinate with appropriate air traffic control to obtain clearance, transit to the targeted location, and accomplish its assigned mission while being responsive to changing parameters such as weather, other aircraft, mission changes, fuel state, and mechanical status. Essentially, an autonomous machine would accomplish the functions currently performed by a pilot operating an RPA. Ideally, the machine would react more quickly, with fewer mistakes, without fatigue, and without any

costly and time-consuming training; it would perform its mission alone, without relying upon a fragile transglobal data link.

Finally, a note on semantics is appropriate. As in any area of rapid technological or scientific advance, the terminology surrounding RPAs continues to evolve. Early studies tended to use different terminology from that used now. Older terms such as "teleoperation" and "unmanned aerial vehicle" (UAV) have been largely replaced with "remotely piloted aircraft." Popular media's wide use of the term "drone" to describe RPAs further confuses the matter. "Drone" is a term the military has long used to refer to aircraft—frequently to converted Vietnam-era F-4 fighter planes—used for airborne target practice. Unless quoting other sources, this study uses the term "RPA" to refer to an aircraft operated by a human pilot not physically present in the aircraft.

Research in machine autonomy is a subset of the rapidly advancing fields of computing and robotics. Recent developments by private corporations and research institutes, particularly those prompted by Defense Advanced Research Projects Agency (DARPA) challenges, have pushed the frontiers outward considerably. The first DARPA Grand Challenge in 2004 required entrants to design and field a completely autonomous vehicle that would traverse a 150-mile course through the Mojave Desert. Successful completion of the challenge would net the winning team a $1 million prize. Of the 15 teams that made it to the starting line, none managed to cover more than 10 miles.[1] However, the next year saw all but one of 23 teams surpass 2004's *best* entrant, and five completed the challenge successfully.[2] Fewer than eight years later, Google researchers made headlines by logging over 300,000 accident-free miles in self-driving vehicles on congested public roads.[3] In July 2013 the United Kingdom joined three US states in legalizing the operation of such vehicles on public roads.[4] This incredibly rapid pace of innovation—driven by the tremendous potential for increased efficiency and safety—clearly demonstrates the trend toward increasing automation and the world-changing implications that entails.

Given that a government research agency sparked much of this progress, it is unsurprising that the US military is this technology's principal beneficiary. Just a few months after the successful completion of DARPA's second grand challenge, the Boeing Phantom Works successfully completed flight tests during which an RPA maintained the contact position behind a KC-135R tanker aircraft, demonstrating the potential for in-flight refueling of autonomous aircraft.[5] More recently, the US Navy successfully landed the remotely piloted X-47B on the deck of the aircraft carrier USS *George H. W. Bush*.[6] A *Foreign Policy* article's coverage of this event clearly conveys the importance of what could otherwise be over-

looked as a simple technical stunt. Gordon Lubold and John Reed observe that "the real news was not the successful launch and now landing of a sleek, jet-powered killer drone on a carrier." Rather, it was that "in the span of a decade, the US military has seen drones transform from primitive, propeller-driven flying lawnmowers to unmanned jets that incorporate all of the features of modern manned aircraft: speed, stealth, high-altitude, sensors, electronic gear, and the ability to carry smart, deadly weapons."[7]

Rise of the Drones

The US military has pursued RPA capabilities for nearly a century, beginning with early experiments during World War I. Thanks to enabling developments in advanced navigation, communications, computing, materials, and satellite communications technologies, however, the past decade has witnessed remarkable growth in both the numbers and use of RPAs. The Department of Defense's (DOD) RPA inventory increased from only 167 aircraft to more than 7,500 from 2002 to 2010.[8] The United States used only a handful of RPAs in the 2003 invasion of Iraq, with just one supporting all of V Corps—the primary US Army combat force. Today, a mission rarely happens without them. Such a trajectory makes it reasonable to postulate future conflicts involving tens of thousands.[9]

Increased usage has led to increased funding—from $667 million in fiscal year (FY) 2001 to more than $3.4 billion in FY 2012.[10] RPAs continue to expand not only in numbers but also in the roles that they fill. Once used almost exclusively for observation roles, RPAs now perform combat operations, search and rescue, and strike coordination and reconnaissance. Suppression of enemy air defenses is a likely near-future role they can fulfill. In fact, almost any function that manned aircraft perform is a potential area for RPA growth. Its development represents a rapidly growing and evolving aspect of the military arsenal—one that shows no signs of slowing in the foreseeable future.

The backbone of the USAF's RPA fleet consists of two airframes manufactured by General Atomics Aeronautical Systems—the MQ-1 Predator and the MQ-9 Reaper. While the US military uses dozens of distinct RPAs for its operations—ranging in size and capability from the backpack-sized AeroVironment Wasp III to the continent-spanning Northrop Grumman RQ-4 Global Hawk—the Predator and Reaper have become the public archetype of the so-called drone. A crew of two operates these aircraft from a remote ground control station (GCS). The GCS and the aircraft are connected by a bidirectional data link that travels via fiber-

optic cable from the GCS to a ground terminal and then via K_u-band radio to a satellite, down to the aircraft, and back again. Just as in manned aviation, the pilot flies the aircraft and maintains responsibility for its overall operation and systems. The sensor operator controls the cameras and lasers contained within the movable turret below the nose. Other than a limited ability to follow a preprogrammed route home if contact with the GCS is lost, the aircraft has essentially no autonomous capability. In fact, it is *less* automated than most manned aircraft.

The principal benefits of this design, aside from greatly reduced risk to the pilot, follow from the removal of human life-support equipment from the plane and a corresponding decrease in weight. Additionally, having no humans in the aircraft makes eliminating many redundant systems acceptable. This weight advantage leads to the increased ability to remain airborne—large RPA loiter times typically hover near 24 hours, 3 to 5 times that of a manned aircraft of similar capability.[11] Those additional hours can be more productive as well. Not only can crews be swapped at any desired interval to maximize alertness but also as many additional experts as required can be accessed simply by sending them the aircraft's video and telemetry feeds.

Finally, those hours aloft come at substantially less cost. A Predator can remain airborne all day on approximately 600 pounds of fuel. In contrast, an RC-135 reconnaissance plane consumes nearly 300,000 pounds while requiring the services of three tanker aircraft—each burning its own fuel and requiring its own crews, maintenance, ramp space, and so forth.[12] Although some disparity exists in the capabilities each provides, the Predator's performance overlaps significantly with that of the RC-135 and has the advantage of a footprint several orders of magnitude less. While a comparison of the Predator with more modest airborne platforms such as the MC-12 is not as dramatic, the concept is the same. RPAs offer a game-changing increase in certain capabilities for a small fraction of the cost. Moreover, their capacities are expanding.

The Predator-style architecture of RPA operation is only one model that the DOD is developing and fielding. Although those systems have stolen much of the media spotlight in recent years, smaller RPAs represent the vast majority of the systems in use at this time. As of FY 2012, the USAF owned a combined total of 215 Predator and Reaper aircraft while the US Army, Navy, and Special Operations Command operated 5,346 RQ-11 Ravens.[13] The Raven, representative of small RPA systems, is a man-portable, battery-powered system launched by hand from a running start—providing ground "personnel with 'over-the-hill' reconnaissance, sniper spotting, and surveillance scouting of intended convoy routes" for up to 90 minutes.[14] These small RPA systems have already

proven exceptionally useful in combat, particularly in light of their modest cost and minimal training requirements. Because their use is confined to short distances, such systems have no need to rely upon satellite communications other than for obtaining precise navigation information, thereby circumventing the bandwidth and latency limitations of their larger peers.

The introduction of weapons to the RPA marked a distinct transition in its evolution. Following successful live-fire tests in 2001, the Predator became the world's first armed RPA.[15] The Reaper followed with a greatly expanded payload and the ability to deliver it twice as quickly. The combination of high-fidelity sensors, extreme loiter time, and precise weapon delivery made these aircraft the first airborne systems to consolidate the entire find, fix, track, target, engage, and assess dynamic targeting chain into a single airborne platform.[16] According to USAF and joint doctrine, this process—known colloquially as the "kill chain"—is how targets are found, prioritized, and pursued. Implementation of the kill chain once required multiple airborne ISR platforms and a dedicated strike asset such as an F-16 or A-10. It can now be accomplished by a single Predator for a fraction of the cost, with no risk to friendly lives, and with less risk of collateral damage.

Contrary to the picture sometimes advanced in the popular press, the Predator's AGM-114 Hellfire missile has become the weapon of choice for ground commanders in Afghanistan due to its precision and comparatively small warhead. Regarding this upward trend in the number and percentage of strikes that RPAs accomplish, the *Los Angeles Times* recently reported that

> the US military launched 506 strikes from unmanned aircraft in Afghanistan last year, according to Pentagon data, a 72% increase from 2011 and a sign that American commanders may begin to rely more heavily on remote-controlled air power to kill Taliban insurgents as they reduce the number of troops on the ground. Though drone strikes represented a fraction of all US air attacks in Afghanistan last year, their use is on the rise even as American troops have pulled back from ground and air operations and pushed Afghan soldiers and police into the lead. In 2011, drone strikes accounted for 5% of US air attacks in Afghanistan; in 2012, the figure rose to 12%.[17]

A 2012 United Nations report indicates that, despite a 72 percent increase in RPA strikes compared to the number that occurred during the previous year, those 506 strikes resulted in 12 percent *fewer* civilian deaths.[18] Pakistan has similar trends, where "RPAs conducted 117 strikes on targets in 2010, up from just 53 in 2009."[19] Additionally, RPA strikes have a disproportionately large effect on the overall war effort because the aircraft—instead of the foot soldiers typically engaged in daily fighting—hunt down and remove specific high-value targets. Such was the case

with the 2013 strike that killed Saudi-born Saeed al-Shihri, the second in command from the Yemen branch of al-Qaeda.[20]

Such disruptive potential has prompted militaries around the world, led by the United States, to aggressively pursue the development of more capable RPA systems. This trend is a result of technological growth and intentional policy. Buried amid the 515-page National Defense Authorization Act of 2001, section 220 states that "it shall be a goal of the Armed Forces to achieve the fielding of unmanned, remotely controlled technology such that (1) by 2010, one-third of the aircraft in the operational deep strike force aircraft fleet are unmanned; and (2) by 2015, one-third of the operational ground combat vehicles are unmanned."[21] Ultimately, this trend is likely to result in fundamental changes to the conduct of war regardless of whether the trend toward remote operation also leads to truly autonomous systems.

Limitations of Current Remotely Piloted Aircraft

Despite their many advantages and promise for the future, RPAs have some shortcomings. While some of these are innate to human operation and shared by manned platforms, others are unique to the technology that enables remote operation. These deficiencies are addressed in two broad categories—technical and human.

The first major obstacle facing current RPAs is finite bandwidth. Measured in bits per second, bandwidth is generally conceptualized as the size of a pipe carrying data. A six millions of bits per second connection can transfer six million bits of digital information in one second. As anyone who has watched a streaming movie has experienced, a higher bandwidth connection equates to a smoother frame rate and higher-resolution video. Contrarily, jerky motion, low resolution, and blocky compression artifacts are symptomatic of low-bandwidth connections. To streaming movie viewers, a finite amount of bandwidth available to their households means that when more than one individual in the household attempts to stream video at the same time, the quality for all must be reduced to fit within that cap.

These limitations are analogous to those of current RPA operations. The largest bandwidth bottleneck occurs in the satellite data link. Due to the substantial expenses involved in the operation of satellite communications systems, a finite amount of bandwidth is available in each satellite's footprint. The capability faces further degradation by environmental factors, such as terrestrial weather and solar particle events. The number of RPAs in simultaneous use has grown, and the resolution of their sensors has evolved much faster than bandwidth capacity. Bandwidth capac-

ity has not kept up with demand, despite the recent launch of the Wideband Global Satellite Communications (SATCOM) system. The system's six satellites, at about $2 billion each, individually provide more bandwidth than the entire Air Force defense SATCOM constellation did for the previous two decades.[22]

Finite bandwidth restricts both the total number of aircraft that can be operated at one time and the amount of data each aircraft can transmit. This data usage is split among three general categories: (1) control information sent to the aircraft from the pilot, (2) telemetry sent from the aircraft back to the pilot (engine performance, fuel status, attitude, altitude, and other similar data), and (3) payload data gathered by the aircraft's sensors.[23] Payload data constitutes the vast majority of overall bandwidth consumption.

For example, the Predator has a fixed-nose camera as well as its more capable sensor cameras. The nose camera is generally not used in flight so that all available bandwidth can be devoted to the more important video feed from the turret sensor. This leaves the crew blind to the area in front of the aircraft, making weather avoidance almost impossible. While some bandwidth can be shared with the nose camera, it is limited to low resolution and low frame-rate updates to minimize impact on the payload full-motion video quality. More than one aircraft has been lost due to this curtailed ability to see and avoid weather.[24]

The problem of limited bandwidth is not unique to RPAs but extends across the entire military. Data usage is growing exponentially for all services—land, sea, and air—and for piloted and remotely piloted assets alike. Satellite communications have become "as vital for unmanned aircraft operations as the fuel that powers their engines."[25] However, unlike manned aircraft where the control information and some of the initial processing occur on board, almost everything current RPAs can do is a result of first getting data from the aircraft to the people who use it. According to a 2012 Congressional Research Service report to Congress, "The finite bandwidth that currently exists for all military aircraft, and the resulting competition for existing bandwidth, may render the expansion of [RPA] applications infeasible and leave many platforms grounded. Ultimately, the requirement for bandwidth grows with every war the United States fights."[26] Thus, a fixed bandwidth caps the number of RPAs that may be simultaneously employed in a theater.

Latency, the second technical limitation of current RPAs, is closely related to bandwidth. Latency is the delay inherent in communications over long distances. Radio waves do not arrive instantaneously although they travel though space at the speed of light. While the light-speed delay is small, the signal processing involved is not. A pilot's control input goes

through many steps before it affects the aircraft on the other side of the world. Analog control positions are converted to digital values, passed through a variety of computers and routers, and then sent via fiber optic cable to a satellite ground station. There, another digital-to-analog translation converts and transmits the positions through the atmosphere to a satellite, which processes and retransmits them back to the aircraft antenna. Positions are once again converted to a digital signal that passes through several individual components within the aircraft, finally moving a small motor connected to a flight control surface. The results of this action are then sent from the aircraft to the pilot via a reverse of the same routing, doubling the amount of time between control input and perceived response. This complex signal path ultimately results in a delay measured on the order of several seconds.[27] While acceptable for some operations, this lag is insufficient for critical actions such as takeoff and landing. It makes in-flight aerial refueling impossible and simple tasks such as voice communication difficult.

The architecture of current RPA systems engenders a third dilemma, that of the vulnerability associated with its communications requirements. Both environmental and malicious sources contribute to this situation. Environmental disturbances such as radio frequency interference, thunderstorms near the satellite ground station, solar flares and charged particle events in space, and poor weather around the aircraft all contribute to reduced bandwidth and, occasionally, complete loss of communications. Improved forecasting and greater power can mitigate some of these problems but cannot completely eliminate them.

Nearly all data transmissions are funneled through only a few nodes, making them a prime target for disruption by enemy computer attack, jamming, or physical destruction. Enemy forces could conceivably attack the satellites directly, which would be more effective than attacking the RPAs themselves. A single attack against the satellite ground terminal could effectively neutralize every airborne RPA in a theater of operations, depriving the United States of a crucial source of information and combat capability.

The final notable technical limitation of current RPA systems relates to the aircraft. Because the aircraft do not carry a live human operator and were initially seen as more or less disposable, their development focused on low cost rather than on reliability. Weapons expert John Pike explains that "for a long time drones were halfway between aircraft and ammo."[28] Aircraft are expected to be reusable, but ammunition is expected to be expended. Attitudes have shifted somewhat in recent years, especially in light of the multimillion-dollar prices of larger RPAs, but the fact remains that Americans tend to value human lives much more

than hardware. RPA systems typically lack the functionality of manned aircraft in which most major systems offer several levels of redundancy to mitigate the consequences of mechanical failures. This deficit frequently places them only a single mechanical failure away from unrecoverable, catastrophic loss.

Likewise, current RPAs lack hardware aimed at enhancing self-preservation. While nearly every manned military and commercial aircraft has a traffic collision avoidance system on board to aid in maintaining safe separation between aircraft, almost no RPAs have one. Also missing are terrain awareness and warning systems and ground proximity warning systems—two common tools used to help prevent controlled flight crashing into the terrain. Because they have evolved in the relatively benign airspace over postinvasion Iraq and Afghanistan, current RPAs are generally unprepared for operation in contested airspace as they are deficient in basic threat warning systems, countermeasures, or stealthy design. Finally, current RPAs have virtually no all-weather capability, possessing neither of the principal weather avoidance tools of their manned counterparts—panoramic windows and color weather radar. This shortfall makes them highly susceptible to operational degradation or loss in even modest weather.[29]

Many of these weaknesses are largely solved in manned aircraft. However, addressing them in the same way as are manned aircraft would negate one of the main advantages of RPAs—a comparatively low cost due to simpler hardware and smaller, lighter airframes.

Human limitations drive the second broad problem area in current RPA operations. While manned aircraft share many of the same problems, some of the aspects unique to RPA operations pose difficulties of their own. Human aircrews are prone to failures of attention, exacerbated by conditions such as fatigue. Lengthy periods of inactivity and monotony punctuated by relatively brief periods of high stress have long been symptomatic of flying operations. Two factors intensify this effect in RPA operations—extremely long missions and the absence of most of the sensory inputs experienced by crews physically present in the aircraft. Both conventional aircrews and RPA operators manifest reduced proficiency stemming from task saturation and "channelized attention" in which key information may be missed. Ultimately, human beings have vulnerabilities that cause their performance in stressful combat and flight operations to fall short of their ideals.

Unfortunately, poor interface design intensifies these problems for RPA crews. Human factors design is an entire branch of research devoted to optimizing interfaces to reduce errors and simplify operations. Pilots widely bemoan the human factors design of RPA cockpits. Information is

frequently buried in tables of numbers and nested computer menus rather than easily accessed through the color-coded gauges and simple physical switches of a traditional aircraft. Col John Dougherty, a Predator operations group commander, addressed this problem at the 2012 Association for Unmanned Vehicle Systems International conference. He stated that "human factors [design] was not integrated into the original design of the Predator. [The engineers] were never given the time, because what was originally a technology demonstration project proved so valuable it was rushed into widespread use. As a result, the percentage of major mishaps caused by human factors is, ironically, higher for Predators than for manned aircraft."[30] The inadequate integration of human factors design into RPA systems was demonstrated clearly in 2010 when "a Predator drone crashed northeast of Kandahar Air Base in Afghanistan after the remote pilot pushed an incorrect button."[31]

The RPA advantage of a long loiter time also creates problems for human operators. Aircraft that can stay aloft for an entire day cause crews to operate in shifts, leading to cumulative fatigue and lifestyle stresses. A recent USAF study reported that the chief stressors for RPA operators correlated with long hours and frequent shift changes.[32] The emotional trauma that results from flying combat operations on a daily basis while continuing a normal civilian lifestyle leads to burnout, clinical distress, post-traumatic stress disorder, and other mental health problems, as well as retention problems for the USAF. For instance, as a Reuters news report indicated, "pressing a button that can lead to someone's death half a world away, then ending your shift to meet family at, say, a child's soccer practice . . . can be difficult for soldiers."[33]

Recognizing the toll this profession takes on its aircrews, the USAF commissioned a 2011 study involving more than 840 Predator, Reaper, and Global Hawk operators. This study found that almost half of those surveyed reported high operational stress (8 or greater on a 10-point scale) and that a quarter suffered from clinical distress; these numbers are similar to those of returning Iraq War veterans.[34] The cumulative results of these human limitations are both minor and life-altering mistakes, slow reaction times, poor decisions, and missed opportunities. Human operators are ultimately the weakest and most error-prone component of the RPA system.

Despite their many benefits, currently fielded RPAs have inherent shortcomings. On the one hand, they are technically constrained by factors including bandwidth scarcity, latency, data-link vulnerability, and a general lack of redundancy and survivability. On the other hand, their human operators impose obstacles on the ability of remotely piloted systems to achieve their potential. Human beings are prone to mistakes in

judgment, inattention, stress, fatigue, and emotional impairment; these problems are at least as detrimental to RPA operations as to manned flying. All of these areas can be addressed, to various degrees, by replacing remote human control with onboard automation.

Automation in Planetary Rovers

Deep Space Exploration and Rover Limitations

NASA is the civilian organization that operates the nation's space program. Although it conducts space operations, its primary function is scientific research. This research falls generally into four categories: aeronautics, human exploration, science, and space technology.[35] Resulting from decades of experience in the design and operation of robotic spacecraft, NASA is one of the foremost research institutions in the world in the fields of robotics and machine autonomy. NASA has successfully landed four mobile rovers on the surface of Mars, beginning with *Sojourner* in 1997 and culminating with the spectacular landing of *Curiosity* in 2012.

In the modern world of ubiquitous wireless data connections, the fact that a picture taken on Mars can be viewed on Earth the same day is easy to take for granted. In reality, data transmission across the vast distances in our solar system is a complex and costly undertaking that relies on the Deep Space Network (DSN). The DSN is a global system composed of three large (230-foot-diameter) steerable parabolic reflector antennas and numerous smaller dishes operated by the Jet Propulsion Laboratory (JPL).[36] These three stations—located in remote parts of the world far from human population centers—are spaced equally around the planet so that an antenna is always pointed toward humanity's far-flung probes. This isolation is necessary to achieve a sufficient signal-to-noise ratio to decipher the original signal. Typical spacecraft transmitters, built to operate within tight power constraints, transmit with a power of about 20 watts—less than one-third the power of a typical household lightbulb. By the time this signal reaches Earth, it "can be as weak as a billionth of a billionth of a watt."[37]

These communications are subject to the same latency and bandwidth limitations as Earth's satellite communications but to an even greater extent. While the speed of light, and thus radio waves, does not contribute significantly to the latency of communications around the earth, it is a real concern over interplanetary distances. For example, round-trip communication delays with Mars vary between around 8 minutes to more than 40 minutes, depending on the relative distance between the planets.

Latency of this magnitude makes direct human control of vehicles across such distances essentially impossible. Bandwidth across such distances, using such low-powered transmitters, is also quite restricted. For a sense of the confines of this communication "pipe," consider that upgrades in 2012 allowed communications with Martian rovers to occur at 256 kilobits per second (kbps); current mobile phones can transmit data at up to 300,000 kbps.[38]

In addition to limitations created by latency and bandwidth between NASA scientists and their distant rovers, a third instance of trouble arose from an unexpected source—increasing rover performance. The *Sojourner* was a suitcase-sized, solar-powered vehicle, while the *Curiosity* is a nuclear-powered vehicle comparable in size to a family minivan. The larger size and greater power offered many advantages, including more scientific instruments, increased speed, and expanded range.[39] The ability to move farther in a given amount of time and collect more data along the way not only created more opportunities for scientific discoveries but also strained the already limited capacity for getting that data back to Earth.[40] The JPL team described the problem:

> Due to advances in rover navigation, traverse ranges are increasing at a rate much faster than communications bandwidth. While the *Sojourner* rover traveled around 100 meters in the entire mission, the drive record for the most distance covered in a single sol (Martian day) is over 220 meters set by the *Opportunity* rover. As this trend in increased mobility continues, the quantity of data that can be returned to Earth per meter traversed is reduced. Thus, much of the terrain the rover observes on a long traverse may never be observed or examined by scientists.[41]

Considering the tremendous cost and risk involved in getting the rovers to Mars, NASA developed the OASIS program to address this trend and improve the quality of scientific data returned. The program allows its rovers to perform many functions without direct human control. These tasks include fundamental navigation, such as slope and obstacle avoidance; power monitoring and management; and basic science work. Because not enough bandwidth is available to transmit every image captured, OASIS determines how best to use its finite resources to maximize science returns. It autonomously decides what to send and what to discard, looks for interesting features along the way, and alters course for closer investigation without waiting for permission from Earth. While these capabilities seem relatively simple in human terms, they represent a challenge to the state of the art in robotics. These changes mark a transition from direct human control toward a truly autonomous machine capable of carrying out its mission and reporting the results back to its human operator.

Conceptual Framework of Autonomy

One of the first tasks in creating the OASIS system was the adoption of a conceptual framework to organize the research. Thomas Sheridan and William Verplank conducted foundational research in this area at the Massachusetts Institute of Technology in 1978. Their detailed study *Human and Computer Control of Undersea Teleoperators* is still referenced frequently in writings concerning machine automation research because its conceptual organization has proven useful over the intervening years. A key concept is that automation is not a binary quality, either present or absent. According to Sheridan and Verplank's model, automation exists along a continuum, which they divide into 10 levels of human-computer decision making (see below table).

Levels of automation in human/computer decision making

1	Human does the whole job, up to the point of turning it over to the computer to implement.
2	Computer helps by determining the options.
3	Computer helps determine options and suggests one, which human need not follow.
4	Computer selects action, and human may or may not do it.
5	Computer selects action and implements it if human approves.
6	Computer selects action and informs human in plenty of time to stop it.
7	Computer does the whole job and necessarily tells human what it did.
8	Computer does the whole job and tells human what it did only if human explicitly asks.
9	Computer does the whole job and tells human what it did; the computer decides if human should be told.
10	Computer does the whole job if it decides it should be done, and, if so, tells human, if it decides human should be told.

Adapted from Thomas B. Sheridan and William L. Verplank, *Human and Computer Control of Undersea Teleoperators*, technical report (Cambridge, MA: Massachusetts Institute of Technology, Department of Mechanical Engineering, Man-Machine Systems Laboratory, 1978), chap. 8, 17–19.

A second piece of the conceptual framework behind the development of OASIS was a model developed by Hui-Min Huang of the National Institute for Standards and Technology. Presented at the Autonomy Levels for Unmanned Systems (ALFUS) working group, this model defines autonomy along a surface in three dimensions: environmental complexity,

mission complexity, and human independence. The goal of the ALFUS working group was creating standardized terminology and metrics to facilitate characterization of remotely piloted systems along a continuum between direct remote control and "full and intelligent autonomy" with "human levels of performance, or better."[42] Within this framework, a cruise missile could be said to have a limited degree of autonomy. It operates without direct human control and exhibits a moderate degree of human independence. However, its rote adherence to a preprogrammed flight path and its ability to interact with the environment with only one response (detonation) means that it achieves a low rating along the environmental and mission complexity dimensions. *Curiosity*, on the other hand, would rank highly along all three axes of this model. It operates in the complex natural environment of the Martian surface and performs a complex scientific and navigation mission with a significant degree of independence from direct human supervision.

A third important concept underpinning autonomy research is the idea of "adjustable autonomy" and its relation to resilience, a key goal of successful autonomous machines. Resilience is the "intrinsic ability of [a human-machine system] to keep or recover to a stable state allowing it to continue operations after a major mishap or in presence of a continuous stress."[43] In more general terms, a resilient system can adapt to dynamic environmental challenges, handle failures, and continue to perform its mission. Imagine a hapless robotic vacuum ineffectually bumping against an unforeseen obstacle until its battery fails, rather than simply going around the obstruction—this exemplifies a lack of resiliency.

Because "the maturity of technology and algorithms does not currently allow designing completely autonomous and sufficiently reliable robots to deal with all possible situations," the interim solution is the principle of adjustable autonomy.[44] This solution helps navigate the gray area between simple remote operation and full autonomy by treating the human-machine relationship as a cooperative system. The machine performs on its own as much as possible but falls back on human assistance when it encounters problems beyond its capacity. Likewise, when human involvement is not immediately available, greater authority for decision making can be given to the autonomous machine. Presumably, as an autonomous system becomes more resilient and trustworthy, it will be granted increasing degrees of autonomy. In the interim, dynamically adjusted cooperation between humans and machines will be the way forward.

The fourth and final concept behind the development of OASIS is the idea of modularity. The modern world would not exist without modularity, a key enabler of the Industrial Revolution. The premise is that large

systems can be broken into components, perfected somewhat independently, and reused elsewhere without being reinvented. In OASIS, each part of the autonomy problem was broken into subcomponents at quite low levels. Each subcomponent was optimized independently and repurposed for other rovers using different hardware. NASA engineers incorporated better algorithms or newly developed sensors, thereby improving the system as a whole without redesigning the entire assembly.

Onboard Autonomous Science Investigation System Components

NASA researchers at the JPL designed a system to endow their rovers with appropriate degrees of autonomy using the four principles as a guiding framework. This effort, dubbed OASIS, is comprised of many subsystems, which the research team categorized into three broad areas—feature detection, data analysis and prioritization, and planning and scheduling.[45] The functional diagram illustrates how these areas are further divided into smaller blocks and share data with one another.
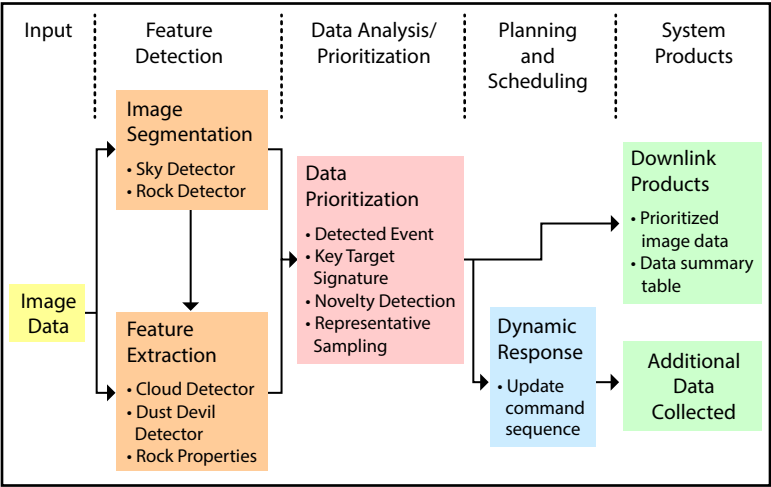
OASIS components and relationships. (*Reproduced from* Rebecca Castano et al., "OASIS: Onboard Autonomous Science Investigation System for Opportunistic Rover Science," *Journal of Field Robotics* 24, no. 5 [May 2007]: 4.)

Two image segmentation modules and three feature extraction modules comprise feature detection. The first image segmentation module, the sky detector, has the function of deciding which portions of a cap-

tured image represent sky and which do not. The OASIS development team tested more than 300 sample images; the chosen algorithms achieved 90 percent accuracy, with errors evenly distributed between false positives and misses.[46] Once the sky detector segmentation is complete, the region of the image determined to be "not sky" is then passed to the second image segmentation module, the rock detector.

The shape, size, texture, location, and distribution of rocks on a planetary surface contain a wealth of scientific information. Researchers developed a variety of algorithms for picking rocks out of image data in the decade leading up to the development of OASIS. These included "stereo-based techniques for finding rocks based on their protrusion from the ground plane, edge-based methods that find closed contours, template-based methods that look for characteristic pixel patterns and methods that detect rocks using their shadows."[47] The OASIS team combined aspects of each, resulting in a system that achieved greater than 90 percent accuracy with few false positives; other algorithms achieved only about 70 percent accuracy.[48]

Once the sky and rock detectors segment the images, the three feature extraction modules perform their tasks on the appropriate image segments. First, the cloud detector determines the number and size of clouds overhead, useful for meteorological studies. This component is more than 93 percent accurate.[49] The dust devil detector tracks and catalogs 85 percent of the dust devils in the rover's vicinity.[50] The most important feature extraction routine is the "rock properties" component, which reliably estimates the albedo (visible light reflectivity), texture, size, and shape of rocks around the rover.[51] Identification of these properties allows the rover to perform the basic functions of a geologist and produce valuable scientific data without relying on a human to process every image.

After the extraction of features, the information moves to the data analysis and prioritization systems. The rover decides which images warrant further action, which ones should be sent back to Earth, and which ones can be discarded. For instance, sky images that contain no clouds need not consume valuable bandwidth. A key part of the processing of rocky images is the idea of "novelty detection." Having 100 images of the same type of rock could be useful, but it is much more likely that novel features contain scientific value.[52] OASIS is able to separate unique rock samples from more common ones by using a combination of different machine-learning algorithms that scientists can optimize. Detection of unique features biases transmission priority and sends a science alert to the next module—the planning and execution system.

The highest-level portion of the OASIS system, the planning and execution module, is designed to run with little interaction with ground controllers. This module receives new science goals from Earth as well as from science alerts generated by the novelty detector. It then modifies the rover command sequence to achieve as many of the goals as possible within resource constraints, such as daylight and electrical power availability. This module executes the rover plan by transmitting commands to the rover's low-level control software and then monitors systems to identify potential problems or opportunities. If the latter was detected, the system controls such situations by adding, moving, or deleting plan activities.[53] Essentially, the planning and execution module is the "brain" of the system and the portion that enables autonomous operation.

The Benefits of OASIS

Development of autonomous capabilities greatly increased the efficiency of NASA's Mars rovers and addressed some of the limitations imposed by the tyranny of distance. First, because the rover can navigate based upon goals set by scientists, an operator does not need to send direct control commands to the rover and wait for confirmation that the commanded movement achieved the desired effect. Latency makes such control exceedingly difficult and inefficient, but goal setting circumvents this problem entirely. Second, autonomy solves much of the bandwidth problem. Because the rovers have some ability to interpret the camera-collected imagery, they can prioritize data for transmission back to Earth. Only images containing novel features or representative samples need to consume valuable bandwidth. Previously, the rover sent a continuous stream of image data as it traversed. The rover can now send a summary report with a few bytes of data equivalent to the following: "I traveled 15 meters today and during the transit passed 18 large rocks and 48 small rocks. All were of the same sedimentary type. No clouds were observed today, but there was one dust devil, and here is a picture of it."

Additionally, the rovers can travel greater distances with less fear of missing meaningful discoveries. Because the novel feature detector collaborates with the planning and execution unit, the rover is capable of making relatively humanlike decisions. Essentially, the rover can reason like this: "I was told to go take a sample from that rock 15 meters away. However, only one-third of the way there, I saw something interesting just 3 meters out of the way. Based on the sun angle and lack of clouds, I've got enough power available to deviate from my primary goal, inspect the new thing, and still achieve my main purpose within my allotted time." As JPL scientists explain, autonomy "enables fast traverse without

sacrificing our understanding of the terrain visited *en route*."[54] Directly micromanaging the rover's operation used to require considerable human oversight. The decreased manpower needed to conduct rover operations and the corresponding reduction in cost are a consequential benefit of the ability to issue high-level goals.

The success of OASIS is most clearly demonstrated in rapidly mounting traverse distance records. *Sojourner*, the first Martian rover, traveled just 0.1 kilometers (km) during its operational lifetime. The rover *Spirit* covered 7.7 km between 2004 and 2010. With OASIS on board, *Opportunity* has covered 35.76 km since 2004—surpassing the *Apollo 15* lunar rover's 1972 record with a human driver.[55] These ever-greater distances are traversed more quickly, easily, and safely and with higher-quality scientific return as a direct result of autonomous operations. The success of NASA's autonomy research has paid tremendous dividends already, and plans are under way to expand the capabilities of these autonomous agents and design more ambitious future missions around those growing capabilities.

Analysis: Application of OASIS to Military RPAs

Conceptual Framework, Trust, and the Modular Approach

NASA's approach to rover automation offers a number of lessons that demonstrate a proven path for military RPA development. The most important of these are not the merits a specific piece of hardware or programming but the overall philosophy that guided its development. This conceptual framework is comprised of two broad areas—an adjustable approach to autonomy and the development of modular systems. Neither of these is a particularly striking observation, but the state of the RPA industry in its current form indicates that emphasizing these lessons is still useful.

Modularity is simply a means for problem solving or engineering in which a larger problem is broken into independent subsections that are attacked one piece at a time. This method was evident throughout OASIS's development. Rather than NASA writing a monolithic "image-processing module," the task of turning camera data into useful scientific and navigation information was broken into multiple parts. Those parts were individually tested, optimized, and then integrated into the whole, creating a larger system with unprecedented capabilities. As new techniques were developed with applicability to one of those modules, a part could be improved without redesigning the entire system. Additionally, the entire OASIS system was designed from the outset to be hardware agnostic,

meaning that it can be used on multiple different rover models. The OA-SIS system was installed on three different rover vehicles (*Rocky 7*, *Rocky 8*, and *FIDO*) during testing before being uploaded to hardware already on Mars.[56] This capability relies on another modular system dubbed the Coupled Layered Architecture for Robotic Autonomy (CLARAty). According to Rebecca Castano et al., CLARAty "is a unified and reusable robotic software architecture that provides a large range of basic robotic functionality and simplifies the integration of new technologies on different robotic platforms."[57]

The thoroughly modular design philosophy achieves numerous benefits. Modularity allows steady progress on individual components performed by different teams that do not have to be at the same location, distributing the workload and allowing competition between different approaches. It encourages greater specialization in each portion of the overall task. From a strictly hardware standpoint, shared components lead to reduced costs, simplified inventory management, and an easier replacement of failed components. The modularity of these designs means that future missions will benefit from the developments already realized without needing to reinvent everything from scratch.

The USAF understands the power of this approach. According to the *United States Air Force Unmanned Aircraft Systems Flight Plan 2009–2047*, "Modular systems with standardized interfaces are required for adaptability, sustainability, and reducing cost."[58] Throughout that document, the service reiterates its intent to foster the creation of a modular RPA platform designed by the defense industry's competitive efforts. NASA's example indicates that this is absolutely the optimum way forward. However, the widely disparate design philosophies, capabilities, and lack of standardization evident throughout the current RPA community indicate this philosophy is not followed. Rather, each airborne system is developed, tested, and fielded independently, each by a different aviation company without sharing components, interfaces, or expertise. This autonomy is representative of the modern military acquisitions process, with the bloated budgets and long development times it connotes. In a clear indication that this "business as usual" approach to RPA development persists, despite the flight plan's recommendations, the US Navy contracted Lockheed Martin, Northrop Grumman, Boeing, and General Atomics to independently design its remotely piloted carrier-launched surveillance and strike aircraft.[59] Each company will design, build, and test its own complete system, despite the significant investment already made in the successful testing of the X-47B for this exact role in 2013.

Aside from the practical benefits of modular design, a larger and less obvious benefit—trust—may well be a deciding factor in the pursuit of

more fully automated RPAs. Just as trust must be earned in human relationships, automated machines must earn the trust of their human operators through proven reliability. According to the *Unmanned Aircraft Systems Flight Plan*, the "trust required for increased autonomy of systems will be developed incrementally."[60] The best way to handle this transition is one component at a time: once each piece of the system is designed and tested exhaustively, operators will be much more likely to rely upon it. Establishing reliability is an important issue for any expensive piece of equipment and dramatically more so when the machine in question is capable of lethal action.

The USAF expects technological advances to allow the creation of a fully autonomous, weaponized RPA within the next decade.[61] Such a weapon would have several appreciable advantages, providing the impetus behind the trend toward increasing automation. However, a responsible state actor would have many complex issues to resolve before fielding such a machine, including legal ramifications, moral questions, acceptability from a social policy perspective, and repercussions for international relations. Some of these obstacles may be insurmountable due to intractable international public perception problems or cultural preferences and will eventually create limits on the technological capabilities. However, those limits may be much more liberal than present-day common sense indicates because of the competitive nature of international relations and military technology development. Granting full autonomy could reduce decision times to the order of microseconds, making human involvement in the process impossible. The first party to remove humans from the decision cycle will have a competitive advantage, creating pressure for other parties to adopt similarly aggressive degrees of autonomy.[62] The USAF defines this method of operation as "human-on-the-loop" to distinguish it from "human-in-the-loop."[63] This challenge is another modern iteration of the military arms race, just as the longbow, the tank, and the space race were in their own times.

While this eventuality seems far afield today, the technology to enable it and the motivation to pursue it are already largely in place. The principal restriction is defined mostly by human trust (or lack thereof) in these systems. NASA's guiding philosophies, including adjustable autonomy and modular design, provide the mechanism for incremental development of that critical relationship.

Feature Detection

Aside from the more philosophical aspects of the development of autonomous RPAs, the feature detection capabilities behind OASIS apply to

the military RPA. In OASIS, feature detection forms the foundation of autonomy. Extracting information from image data is how the machine "knows" what's going on around it. Knowledge of the surrounding environment—coupled with defined goals and other sources of information—forms the basis of independent reasoning. On a Mars rover, the image segmentation algorithm and the sky, rock, cloud, and dust devil detectors recognized features of interest. A number of analogous routines would not only form the basis of an automated RPA but also provide immediate benefits to human operators in a cooperative relationship short of full autonomy.

Similar to the cloud and dust devil detector components of OASIS, a cloud detection capability would be useful to a system such as the Predator. Because these aircraft currently have limited all-weather capability, avoidance of clouds is a crucial part of safe operations. That task is particularly difficult for a pilot who primarily views the area through a telescopic camera that must be pointed at the ground for the aircraft to perform its mission. This extremely close view is often referred to as a "soda straw" perspective by RPA pilots, referring to the almost complete lack of situational awareness about the aircraft's surroundings.

As previously described, the Predator is equipped with a nose camera that provides much lower fidelity than the main turret-mounted camera systems, and it must take precious bandwidth from those cameras if it is to be used at all. A reliable onboard weather detection system could use the nose camera full-time with no satellite communication penalties. In fact, that camera could be much higher resolution or augmented by cameras providing a full 360° view. The pilot would need to be sent only the range and bearing to the nearest cloud and an assessment of the threat it presents. Such a system would undoubtedly lead to greater operational safety, fewer lost airframes, and greater freedom for the human aircrew to focus on operational tasks.

A second useful feature detection capability would be a "human detector." Warfare is fundamentally a human endeavor, and keeping track of people is an essential task—particularly for the counterinsurgency environment in which current RPAs have evolved. The ability to automatically detect human shapes within an image, track them as they move about, and classify them into broad categories such as male, female, or child would be exceptionally useful; this activity currently accounts for a colossal investment of human effort. A reliable automatic human detector would largely obviate the need to send full-motion video back to human analysts, thereby nearly eliminating the bandwidth problem. Rather than sending volumes of video data, the aircraft could relay simple informational summaries, including the number of people, arrivals, depar-

tures, and still images. This feature would not only save human time, money, and bandwidth but also protect lives and US national interests.

If such a system seems implausibly futuristic, then it will come as quite a surprise to learn that it already exists. Sentient, a company that develops computer vision and artificial intelligence software for defense and civilian applications, markets such a system under the name Kestrel Land Moving Target Indicator. This system takes the video feed from an airborne ISR platform and automatically highlights moving figures—even those too small or slow moving for a human to reliably detect. It can even create a visible track behind moving objects, allowing the motion of people and vehicles to be accurately predicted and making those positions obvious even when the movement stops or becomes obscured by vegetation.[64] While human operators can usually perform some of these tasks when focusing on a single object, no one can reliably perform this task with a screen filled with individuals all moving independently. Machine performance of this somewhat complex image-analysis task already surpasses human ability. Augmenting human skills with advanced automation leads to increased reliability, safety, and capability.

Despite their best efforts, human operators are systematically prone to errors. One such error is channelized attention. Especially in stressful situations, people tend to focus on only a few of their tasks to the almost complete exclusion of other inputs. During a strike, an RPA pilot is focused on weapons employment checklists, weapons engagement zone calculations, aircraft positioning, communication with other assets, and compliance with instructions from the joint terminal attack controller. The presence of an innocent bystander on the corner of the screen is an easily missed detail but one of utmost importance. So-called collateral damage does substantial harm to the larger war effort, the stature of the United States, and the psyche of the people who make such errors. Highlighting potential civilians within the risk-estimate distance of the selected ordnance with an automated human shape recognition module would be relatively trivial but would provide tremendous benefits.

Similar to human shape detection, automated detection of enemy military hardware—such as ground vehicles, antiaircraft weapons, and other aircraft—would be extremely useful. Such capability would provide the basis of a threat recognition capability, especially when combined with other types of sensors and information from other platforms. This awareness of military threats could be combined with the cloud detection module to create an overall threat avoidance system. Doing so would help solve one of the most severe limitations of current RPA systems— the reliance on permissive airspace, won through air superiority and earned by manned assets at considerable human risk. Outside the theater

of war, a threat avoidance system would provide a means of safely maintaining separation from other aircraft, a major obstacle to increased RPA use within the national airspace system.

These programmed routines would form the foundation of a more autonomous RPA. However, before a fully independent system uses that information, it will provide an extremely valuable source of additional data to a human operator. In this role as a decision aid for humans, such feature recognition systems may be refined and perfected, thereby earning the trust necessary to allow them to perform more indepently.

The Brains: Analyzing, Prioritizing, Scheduling, Planning, and Goal Setting

Just as NASA rover operation evolved from simple remote control to nearly full autonomy, such is the likely path toward automation in military RPAs. Progress over the next several years should focus on headway in two goals: creating reliable automated decision aids and making individual subsystems fully automated but still controlled by a human operator. The aforementioned feature detection capabilities are a substantial part of the first goal. Fuel management in a current-generation Predator is an example of the second.

The aircraft's fuel system reports the quantity of remaining fuel to the pilot. Although helpful information, it is simply a number. The pilot and others really want to know the amount of time the aircraft can remain on station. This determination requires some simple math; the pilot estimates the distance to the recovery base using a computerized chart program, the average fuel flow during the transit based on past averages, and the expected ground speed with an estimated wind factor and subtracts the required reserve fuel at landing. The pilot uses that data to determine the amount of time the aircraft may remain overhead its objective. While not complex, this process takes mental effort and involves a good deal of estimation, leading to inexact estimates. An automated system would use exact numbers, all available to the aircraft, to provide a "remaining station time" directly. This improvement would prevent errors, reduce human workload, and maximize time on station without lowering margins for safety. Similar but more dramatic improvements are easily attainable with other systems—especially weapons delivery.

Ultimately, the human operator should be free to make higher-level decisions in the form of goals that the aircraft pursues on its own, reporting progress and relevant status when required. This strategy allows for the development of trust in individual mission subsets and represents an adjustable level of autonomy that paves the way for more complete au-

tonomy. Eventually, RPAs are likely to be operated similar to current NASA rovers—operators establish high-level goals that the autonomous machine accomplishes by planning and scheduling activities, monitoring resources and other constraints, and observing limitations imposed by the operators. Environmental changes, such as threats or weather, cause the system to respond appropriately by making changes to current goals and shuffling plans accordingly.

Projected Benefits

Like most complex subjects, the path to autonomy will probably not be marked by a single definite threshold. Rather, the evolution of these machines will take place in small increments, with varying benefits becoming evident over time. Initially, the expediency of RPA automation will lead to improvements similar to those achieved by NASA's rovers. Local processing will result in a reduced need to transmit everything to the operator. This function will slow the rapidly growing demand for more bandwidth, allowing greater numbers of RPAs to operate within the same bandwidth capacity. As RPAs become more autonomous, they will be less dependent upon their direct link to an operator, limiting the vulnerability created by funneling so much communication through a single ground facility. Autonomous action will also overcome latency issues to reduce reaction times substantially. RPAs could survive and operate in nonpermissive airspace, taking on roles currently performed by manned assets at great risk to human life. It will also be possible to increase the number of vehicles a single human may command, increasing the efficiency of the entire operation, reducing costs and training requirements, and expanding military capability. Integration of automated decision aids will reduce accident rates, decrease human stress, and limit the occurrence of collateral damage in kinetic strikes. While delays and missteps will undoubtedly occur along the way, NASA has already demonstrated that an appropriate mixture of theory and technology can achieve significant benefits through the development of autonomous machines.

Conclusion and Recommendations

Recommended Research and Development Road Map

Several specific lessons from NASA's OASIS research apply directly to the development of military RPAs. First, the USAF must employ an appropriate overall philosophy to this initiative based on the concepts of modular design and adjustable autonomy. Modular design creates a system of steady advancement, shared components, and competitive devel-

opment at a subsystem level. The principle of adjustable autonomy allows for reliance on human oversight and autonomous behavior according to each situation's demands. Robust feature detection capabilities are a good place to start, initially as decision aids to human operators and then as the basis for onboard autonomous decision making. At the same time, more complete automation of aircraft systems will lead to enhanced safety and pave the way for full autonomy. As the reliability of various components improves, human roles will transition from direct remote control to oversight. Individual phases of flight such as takeoff, landing, and aerial refueling will eventually become autonomous. At that point, the leap to fully autonomous operation will be a relatively small distinction, leaving the human operators in a strictly supervisory role.

Potential Pitfalls

The most likely points of failure along this path to autonomy are not technical but within the domain of human acceptance. The ethical, legal, and cultural implications of this transition are more difficult to navigate than the technical challenges. The public perception of current systems is already substantially negative; popular culture is rife with apocalyptic scenarios in which lethal machines escape human control.[65] Aside from the general suspicion and lack of trust with which people view automated systems are difficult legal questions. Who is held accountable when an autonomous RPA kills the wrong person? Is it the combatant commander who authorized its use, the engineer who built it, the programmer who wrote its code, or some other entity? Is it ethical to allow such a thing to happen in the first place? Furthermore, what happens to the concept of *jus ad bellum*, how one determines the justice of going to war, when the assumed risk of going to war involves not human life but simple hardware? Will political leaders have the same restraint when contemplating the authorization of military force? These discussions may test the very foundations of the just war tradition that has evolved through centuries of warfare.

While a full discussion of these issues is well beyond the scope of this study, there are several reasons to believe they will be resolved as a matter of course. First, automated decision aids are already employed in military operations, and autonomous weapon systems are already in use—the defensive antimissile guns on US Navy ships are an example. The Phalanx Close-In Weapons System uses autonomous radar-guided, 20-millimeter Gatling guns that react faster than human operators to defend against antiship missiles.[66] Second, it is important to recognize that the standard against which the performance of an autonomous machine must be com-

pared is not perfection but human performance. The military allows 18-year-old Soldiers to wield lethal weapons and grants them the authority to use those weapons in extremely stressful situations. Human performance is sometimes surprisingly good, but it is far from perfect—witness the high-profile death of Cpl Pat Tillman as a result of friendly fire.[67] While every human being, despite military training's best efforts, is an untested individual, machines offer the opportunity for standardized performance and steady improvement. Machines sidestep many of the sources of error humans in combat are susceptible to—fear, fatigue, anger, prejudice, and many others.

Finally, these issues need not await technological developments for us to seek solutions. In fact, investigations into these concerns are already under way. The *United States Air Force Unmanned Aircraft Systems Flight Plan 2009–2047* recognizes the importance of these issues and outlines a clear transition from "man-in-the-loop" to "man-on-the-loop" operations.[68] A 2008 US Navy study, *Autonomous Military Robotics: Risk, Ethics, and Design*, discusses these topics and others in some detail. "The Implications of Drones on the Just War Tradition" addresses the legal and moral aspects of a nation's use of force when it does not risk its own Soldiers' lives. "Embracing Autonomy" is an excellent *Air and Space Power Journal* article describing the justifications and challenges associated with this transition.[69] The popular media has broached these subjects as well. Ultimately, many of the challenges that lie ahead on the path to autonomous aerial war machines are neither unknown nor insurmountable.

A Robotic Future

Despite the science fiction quality of this subject, the United States will follow the path to autonomous RPAs. The likely benefits are too great, the challenges are manageable, and the risk of a competitor perfecting this technology first is unpalatable. However, the USAF must be deliberate in its pursuit of this goal to ensure that its implementation is efficient and proceeds along lines that comply with our cultural values. Moreover, it must seek the development of technology that furthers human goals. NASA's long expertise in the development of machine autonomy demonstrates the principles and technological features that make the goal of autonomous RPAs attainable.

Notes

(All notes appear in shortened form. For full details, see the appropriate entry in the bibliography.)

1. Walton, "Robots Fail."

2. Iagnemma and Buehler, "Journal of Field Robotics."

3. Rosen, "Google's Self-Driving Cars."

4. Farivar, "Government Approves Autonomous Cars."

5. Doug Kinneard, "Boeing, U.S. Air Force Demonstrate UAV Automated Aerial Refueling Capability," Boeing press release, 11 December 2006, http://www.boeing.com/news/releases/.

6. Neela Debnath, "Unmanned X-47B Jet Makes History after Successfully Landing on Aircraft Carrier," *Independent*, 11 July 2013, http://www.independent.co.uk/news/world/.

7. Lubold and Reed, "This Drone."

8. Gertler, *U.S. Unmanned Aerial Systems*, 8.

9. Singer, "Robots at War," 30–48.

10. Gertler, *U.S. Unmanned Aerial Systems*, 8.

11. Ibid.

12. Ibid.

13. Ibid.

14. Ibid., 45.

15. Lee, "Embracing Autonomy," 78.

16. Joint Publication 3-60, *Joint Targeting*, II-12.

17. Shashank Bengali and David S. Cloud, "U.S. Drone Strikes Up Sharply in Afghanistan," *Los Angeles Times*, 21 February 2013, http://articles.latimes.com/2013/feb/21/world/la-fg-afghanistan-drones-20130222.

18. United Nations Assistance Commission in Afghanistan, *Afghanistan Annual Report 2012*, 1.

19. Lee, "Embracing Autonomy," 86.

20. Ahmed Al-Haj, "Saeed Al-Shihri Dead."

21. National Defense Authorization Act of 2001, sec. 220 (a).

22. Covault, "UAVs Drive SATCOM Modernization."

23. Valdez, "Predator UAV Works."

24. Nelson, "Just a Weather Forecast."

25. Covault, "UAVs Drive SATCOM Modernization."

26. Gertler, *U.S. Unmanned Aerial Systems*, 17.

27. Valdez, "Predator UAV Works."

28. Rizzo, "Raining Drones?"

29. Nelson, "Just a Weather Forecast."

30. Quoted in Freedberg, "Too Many Screens."

31. Rizzo, "Raining Drones?"

32. Elizabeth Bumiller, "Air Force Drone Operators Report High Levels of Stress," *New York Times*, 18 December 2011.

33. Stewart, "Overstretched U.S. Drone Pilots."

34. Bumiller, "Air Force Drone Operators."

35. NASA, "About NASA."

36. NASA, *Deep Space Network*.

37. Ibid.

38. Ibid.

39. Bentley, "Space Probes," 42–45.

40. Castano et al., "OASIS," 2.

41. Ibid.

42. Huang, "Autonomy Levels."

43. Zieba et al., "Principles of Adjustable Autonomy," 194.

44. Ibid., 193.

45. Castano et al., "OASIS," 382.

46. Ibid., 383.

47. Ibid.

48. Ibid.

49. Ibid.

50. Ibid., 384.

51. Ibid.

52. Ibid., 385.

53. Ibid., 387.

54. Thompson et al., "Onboard Science Data Analysis," 2.

55. Kremer, "Opportunity Mars Rover."

56. Castano et al., "OASIS," 390.

57. Ibid.

58. *Aircraft Systems Flight Plan*, 14.

59. Gallagher, "Navy Commissions New Designs."

60. *Aircraft Systems Flight Plan*, 41.

61. Ibid., 50.

62. Ibid., 41.

63. Ibid.

64. Sentient, "Kestrel Land MTI."

65. Singer, "Robots at War," 30–48.

66. US Navy Fact File, "MK-15 Phalanx."

67. "What Really Happened?"

68. *Aircraft Systems Flight Plan*, 42.

69. Lin, Bekey, and Abney, "Autonomous Military Robotics"; Brunstetter and Braun, "Implications of Drones"; and Lee, "Embracing Autonomy."

Abbreviations

ALFUS   Autonomy Levels for Unmanned Systems
CLARAty   Coupled Layered Architecture for Robotic Autonomy
DARPA   Defense Advanced Research Projects Agency
DOD   Department of Defense
DSN   Deep Space Network
FY   fiscal year
GCS   ground control station
ISR   intelligence, surveillance, and reconnaissance
JPL   Jet Propulsion Laboratory
kbps   kilobits per second
km   kilometer
NASA   National Aeronautics and Space Administration
OASIS   onboard autonomous science investigation system
RPA   remotely piloted aircraft
SATCOM   satellite communications
UAV   unmanned aerial vehicle

Bibliography

Ahmed Al-Haj. "Saeed Al-Shihri Dead: Al Qaeda Branch Number 2 Confirmed as Killed in Drone Strike." Associated Press, 17 July 2013.

Bentley, Molly. "Space Probes on Their Own." *New Scientist* 203, no. 2725 (September 2009): 42–45.

Brunstetter, Daniel, and Megan Braun. "The Implications of Drones on the Just War Tradition." *Ethics and International Affairs* 25, no. 3 (Fall 2011): 337–58.

Castano, Rebecca, Tara Estlin, Robert C. Anderson, Daniel M. Gaines, Andres Castano, Benjamin Bornstein, Caroline Chouinard, and Michele Judd. "OASIS: Onboard Autonomous Science Investigation System for Opportunistic Rover Science." *Journal of Field Robotics* 24, no. 5 (May 2007): 379–97.

Covault, Craig. "UAVs Drive SATCOM Modernization." *Defense Media Network*, 26 October 2010. http://www.defensemedianetwork.com/stories/uavs-drive-satcom-modernization/.

Farivar, Cyrus. "UK Government Approves Autonomous Cars on Public Roads before Year's End." *ArsTechnica*, 16 July 2013. http://arstechnica.com/tech-policy/2013/07/uk-govt-approves-autonomous-cars-on-public-roads-before-years-end/.

Freedberg, Sydney J. "Too Many Screens: Why Drones Are So Hard to Fly, So Easy to Crash." *Breaking Defense*, 7 August 2012, http://breakingdefense.com/2012/08/07/too-many-screens-why-drones-are-so-hard-to-fly-and-so-easy/.

Gallagher, Sean. "Navy Commissions New Designs on Armed, Carrier-Launched Drones." *ArsTechnica*, 16 August 2012. http://arstechnica.com/tech-policy/2013/08/navy-comissions-new-designs-on-armed-carrier-launched-drones/.

Gertler, Jeremiah. *U.S. Unmanned Aerial Systems*. Washington, DC: Congressional Research Service, 3 January 2012.

Goth, Gregory. "Software on Mars." *Communications of the Association for Computing Machinery* 55, no. 11 (November 2012): 13–15.

Huang, Hui-Min. "Autonomy Levels for Unmanned Systems (ALFUS)." National Institute of Standards of Technology. ALFUS Working Group, 2008.

Iagnemma, Karl, and Martin Buehler. "Editorial for Journal of Field Robotics—Special Issue on the DARPA Grand Challenge." *Journal of Field Robotics* 23, no. 8 (2006): 461–62.

Joint Publication 3-60. *Joint Targeting*, 13 April 2007.

Kremer, Ken. "Opportunity Mars Rover Blazes Past 40 Year Old Space Driving Record." *Universe Today*, 18 May 2013. http://www.univer-

setoday.com/102212/opportunity-mars-rover-blazes-past-40-year-old-space-driving-record/.

Lee, Caitlin H. "Embracing Autonomy: The Key to Developing a New Generation of Remotely Piloted Aircraft for Operations in Contested Air Environments." *Air and Space Power Journal* 25, no. 4 (Winter 2011): 76–88.

Lin, Patrick, George Bekey, and Keith Abney. "Autonomous Military Robotics: Risk, Ethics, and Design." *Philosophy,* 20 December 2008, 1–112.

Lubold, Gordon, and John Reed. "This Drone Just Made History. But Now Where Does It Go?" *Foreign Policy*, 10 July 2013. http://foreignpolicy.com/2013/07/10/this-drone-just-made-history-but-now-where-does-it-go/.

National Aeronautics and Space Administration. "About NASA: What Does NASA Do?," November 2014. http://www.nasa.gov/about/highlights/what_does_nasa_do.html.

———. *Deep Space Network*. Brochure, 2005.

National Defense Authorization Act for Fiscal Year 2001. Pub. L. 106–398. 114 Stat. 1654A-38, sec. 220 (a). 106th Cong., 2nd sess., 30 October 2000.

Nelson, Lee. "More than Just a Weather Forecast: The Critical Role of Accurate Weather Data in UAV Missions." *Defense Update*, June 2009. http://defense-update.com/features/2009/june/weather_and_uav_operations.html.

Rizzo, Jennifer. "Raining Drones?" *CNN*, 30 December 2011. http://security.blogs.cnn.com/2011/12/30/raining-drones/.

Rosen, Rebecca J. "Google's Self-Driving Cars: 300,000 Miles Logged, Not a Single Accident under Computer Control." *Atlantic*, 9 August 2012. http://www.theatlantic.com/technology/archive/2012/08/.

Sentient. "Kestrel Land MTI: Automated Moving Target Indication." Accessed 2 October 2015. http://www.sentientvision.com/defence_products_airborne_kestrel_land_mti.html.

Sheridan, Thomas B., and William L. Verplank. *Human and Computer Control of Undersea Teleoperators*. Technical report. Cambridge, MA: Massachusetts Institute of Technology, Department of Mechanical Engineering, Man-Machine Systems Laboratory, 1978.

Singer, P. W. "Robots at War: The New Battlefield." *Wilson Quarterly* 33, no. 1 (Winter 2009): 30–48.

Stewart, Phil. "Overstretched U.S. Drone Pilots Face Stress Risk." Reuters, 19 December 2011.

Sullins, John P. "RoboWarfare: Can Robots Be More Ethical than Humans on the Battlefield?" *Ethics and Information Technology* 12, no. 3 (September 2010): 263–75.

Thompson, David R., Robert C. Anderson, Benjamin Bornstein, Nathalie A. Cabrol, Steve Chien, Tara Estlin, and Terry Fong et al. "Onboard Science Data Analysis: Implications for Future Missions." Planetary Science Decadal Survey Review Paper. Pasadena, CA: Jet Propulsion Laboratory, California Institute of Technology, 15 September 2009.

United Nations Assistance Commission in Afghanistan (UNAMA) and Office of the High Commissioner for Human Rights (OHCHR). *Afghanistan Annual Report 2012: Protection of Civilians in Armed Conflict*. Kabul: UNAMA, OHCHR, February 2013.

United States Air Force. *United States Air Force Unmanned Aircraft Systems Flight Plan 2009–2047*. Washington, DC: Headquarters Air Force, 18 May 2009.

United States Navy Fact File. "MK-15 Phalanx Close-In Weapons System," 15 November 2013.

Valdez, Robert. "How the Predator UAV Works." HowStuffWorks.com. http://science.howstuffworks.com/predator6.htm.

Walton, Marsha. "Robots Fail to Complete Grand Challenge: $1 Million Prize Goes Unclaimed." *CNN*, 6 May 2004. http://www.cnn.com/2004/TECH/ptech/03/14/darpa.race/.

"What Really Happened to Pat Tillman?" *CBS News*, 11 February 2009. http://www.cbsnews.com/news/what-really-happened-to-pat-tillman/.

Zieba, Stéphane, Philippe Polet, Frédéric Vanderhaegen, and Serge Debernard. "Principles of Adjustable Autonomy: A Framework for Resilient Human-Machine Cooperation." *Journal of Cognition, Technology and Work* 12, no. 3 (September 2010): 193–203.